

Presented at ISER'00, Honolulu, 2000.

Ladar-based Discrimination of Grass from Obstacles for Autonomous Navigation

Jose Macedo
University of San Diego
Industrial and Systems Engineering
San Diego, CA 92110
jmacedo@ucsd.edu

Roberto Manduchi, Larry Matthies
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA 91109
{manduchi,lhm}@telerobotics.jpl.nasa.gov

1. Introduction

Abstract: Autonomous navigation in vegetated terrain requires the ability to discriminate obstacles from grass, a non-trivial problem when the sensorial world of the robot is based only on range information as provided, for example, by a laser rangefinder (ladar). We present a statistical analysis of the range data produced by a single-axis ladar in different situations, including the case of an obstacle partially occluded by grass. Such analysis inspired a simple classification algorithm, which has been tested on real range data acquired by JPL's urban robot.

2. Introduction

Autonomous vehicles have great promise for applications in the military, agriculture, space exploration, and other domains. Moreover, rapid progress in miniaturization and improved cost-effectiveness of navigation sensors, cameras, and computers is accelerating the maturation of robotic vehicles. However, a key limitation remains for domains in which robots must navigate in tall grass, small bushes, or forested areas, because existing perception systems cannot do effective obstacle detection in these situations. Most obstacle detection systems to date rely exclusively on range data from ladar, stereo vision, radar, or ultrasonic sensors to perceive scene geometry and assume implicitly that the scene consists of relatively large, solid surfaces [7]. When driving in vegetated terrain, the notion of "obstacle" needs to be revisited. For example, a small bush can be considered an obstacle based solely on geometric speculation, although it probably can be driven over without damaging the vehicle. Thus, for efficient navigation in vegetated terrain, a higher level of reasoning must intervene, based on both the geometric description of the scene and the composition of the terrain cover.

In this paper, we are interested in determining whether an "obstacle" is a rock (non-traversable) or a patch of grass (traversable). Terrain cover classification can be based on color features [1], but such an approach won't work at

night. Visual texture is another promising approach, but it is computationally expensive and the technology is not mature yet. In this work, we discuss a simple approach based on the statistical analysis of the range data as provided by a laser rangefinder (ladar). Intuitively, range data on grass and bushes will be spatially scattered, while range data on bare soil or rock will tend to be more “regular” and lie on a relatively smooth surface.

In the following section we derive theoretical probability distributions for range data from a single-axis ladar in vegetated terrain. We consider a number of “canonical” situations, including the case of a field of randomly distributed grass containing a partially occluded rock. The theoretical results of Section 3 are validated in Section 4, where we show histograms of real range data. We also introduce a simple and fast algorithm for the classification of grass versus rocks based on statistical measures over moving windows. The range data used in the experiments was collected by the urban robot developed by JPL as part of the Tactical Mobile Robotics (TMR) Program funded by DARPA (see Figures 4,8). This robot is equipped with stereo cameras, an omnidirectional camera, an uncooled thermal infrared camera, and a 2-axis scanning ladar (although the data for this work has been acquired by rotation around the vertical axis only) [2, 6]. Autonomous navigation capabilities to date include obstacle avoidance, visual servoing to goals, and autonomous stair climbing.

3. Statistics of Range Data on Grass

The statistics of range measurements can provide us with precious information about the terrain cover. We introduce here a model for the range distribution which can be used to design a classifier of grass versus other obstacles, as discussed in Section 4. It is assumed here that the laser rotates around a vertical axis, and that the laser beam width is infinitesimal. We also assume that the ladar always receives a return when the laser ray hits a surface, and that the measurements are noiseless. These assumptions are discussed in more detail in Section 3.4

In the next subsections, we will consider three “canonical” situations: when the robot is in the middle of an homogeneous field of grass, when the robot is placed at a certain distance from a patch of grass, and when a rock is partly visible through the grass.

3.1. Case 1: Homogeneous grass field

To model an homogeneous field of grass, we will assume that the blades of grass have constant circular section of diameter d , and that the centers of the grass blades are distributed according to an isotropic Poisson point process in space with intensity λ [4]. This means that within a unit area of soil we expect to find λ blades¹. In the case of an infinitesimal laser ray, one easily proves that the probability density function (pdf) of the range r is

$$p(r) = \lambda d e^{-\lambda d r} U(r) \quad (1)$$

¹This model is not entirely correct because it assumes that grass blades can intersect [5].

where $U(\cdot)$ is the Heavyside function. In other words, we expect the range data on grass to behave as an exponential distribution with mean $\mu = 1/\lambda d$. A similar exponential behavior was predicted and observed in the case of range data measured in a forest [3], as a consequence of the random distribution of tree trunks.

A synthetic example of grass distribution is shown in Figure 1. Each grass blade has diameter d equal to 20 mm, and there are on average $\lambda=500$ blades per square meter. It is assumed that the ladar acquires data with an angular period of 0.5° . The blades of grass that are hit by the laser ray are represented with a red kernel. Figure 1 also shows the normalized histogram of the range measured over 500 trials. An exponential density curve has been fitted to the data and superimposed to the histogram in the figure. Such best-fitting exponential has mean μ equal to 0.098/m, which is very close to the expected value of 0.1/m.

From the pdf of the range (1) we may compute its second-order moment (variance, σ^2) and its third-order moment (skewness², sk), which will be used in Section 4:

$$\sigma^2 = \mu^2 = 1/(\lambda d)^2, \quad sk = 2\mu^3 = 2/(\lambda d)^3 \quad (2)$$

The skewness measures the degree of symmetry of a distribution around its mean. Positive (negative) values of the skewness indicate that the distribution extends towards the right (left) tail. The case $sk = 0$ corresponds to a symmetric distribution, such as the normal. In the case under exam the skewness is positive, meaning that the distribution is skewed to the right.

3.2. Case 2: Grassy patch seen from the distance

An instance of this case is shown in Figure 2 ($\lambda = 85/\text{m}^2$, $d = 30$ mm). The robot is at a distance $D = 200$ mm from the rectilinear edge of a patch of grass. The distribution of the range data along a fixed line oriented at an angle α with respect to the normal to the patch edge is now an exponential shifted by a value of $D/\cos(\alpha)$. However, the pdf of the range computed over the whole angular span (equal to 90° in this example) does not have an exponential shape. Indeed, the normalized histogram shown in Figure 2 (computed over 500 trials) shows a heavier tail than in the case of an exponential density.

3.3. Case 3: Rock behind the grass

Suppose the robot is looking at a patch of grass which contains a rock (or any non-traversable object). Clearly, if the rock is right in front of the robot, the distribution of the range concentrates around the actual distance of the rock. If the rock is very far away from the robot, all the rays from the ladar are likely to be stopped by the grass before reaching the rock surface. However, if the rock is at an intermediate distance, it is quite possible that some rays from the ladar reach the rock surface, while the other rays hit some blades of grass on the way. This situation is shown in Figure 3, where the parameters of the

²By identifying the skewness sk with the third-order moment of the distribution, we adhere to the definition given in [8]. Note that other authors (e.g. [4]) define the skewness as the third-order moment divided by $\sigma^3/2$.

grass distribution are $\lambda = 300/\text{m}^2$, $d = 10 \text{ mm}$ (the rock surface is represented by the blue circles.) To simplify the computation of the range distribution, we will assume that all points of the rock surface are at the same distance T from the ladar (this hypothesis is clearly not verified in Figure 3, where the rock surface is at a slanted angle). In this case, the pdf of the range r is simply an exponential truncated at $r = T$, followed by a peak at $r = T$ of area $e^{-\lambda d T}$. The larger the distance T to the rock, the smaller the area of the peak, and the further to the right its position. This expected behavior can be noted in the histogram of Figure 3 (computed over 3000 trials within the angular sector subtended by the rock surface). The histogram is clearly bimodal, although the second peak spreads out due to the slant of the rock surface. The mean μ , variance σ^2 and skewness sk of the range are easily computed from our model distribution:

$$\mu = \frac{1}{\lambda d} (1 - e^{-\lambda d T}) \quad (3)$$

$$\sigma^2 = \frac{1}{(\lambda d)^2} - e^{-\lambda d T} \frac{2T}{\lambda d} + \frac{e^{-2\lambda d T}}{(\lambda d)^2} \quad (4)$$

$$sk = \frac{2}{(\lambda d)^3} - e^{-\lambda d T} \frac{3T^2}{\lambda d} - e^{-2\lambda d T} \frac{6T}{(\lambda d)^2} - e^{-3\lambda d T} \frac{2}{(\lambda d)^3} \quad (5)$$

By comparing (5) with (2), one maintains that the skewness in this case is always smaller than in the case of homogeneous grass, and indeed it takes on negative values for sufficiently small T . In other words, the presence of the peak centered at $r = T$ makes the distribution more skewed to the left.

3.4. Validity of our theoretical assumptions

While the theoretical results derived in the previous subsections are useful to understand the behavior of the range data, and indeed have inspired the simple classification algorithm of Section 4, we should comment on the shortcomings of our analysis. Firstly, the assumption that the ladar always receives a return when the laser ray hits a surface is not very realistic. For example, it is quite possible that the ladar does not read a return when the ray hits the outer edge of a grass blade. Secondly, the return is integrated over a non-infinitesimal time window while the ladar is revolving. This means that the measured range actually represents the average value of the actual profile within a small angular sector. The measured range is thus a “smoothed” version of the actual range profile. Finally, a more realistic analysis would take into account the fact that laser ray has a non-null divergence γ (e.g., the laser used in the experiments has $\gamma = 1 \text{ mr}$). It can be shown that in this case the distribution of the range along a fixed direction can be modeled by the following form:

$$p(r) = \lambda d(1 + ar)e^{-\lambda d(r-D)(1+a(r+D)^2/2)}U(r-D) \quad (6)$$

where D is the distance to the edge of the grassy patch along the ray and $a = 2 \tan(\gamma/2)/d$. For small values of γ and D , the distribution of the range data does not differ substantially from the infinitesimal ray case. This is because the ray will normally hit a blade of grass before “thickening”. However, if the

patch of grass is far from the ladar (i.e., for large enough values of aD^2), it is seen from (6) that the effect of finite ray size cannot be neglected. In this case, the ray will have thickened noticeably before hitting a blade of grass. Intuitively, this corresponds to using an infinitesimal ray in a field of grass where the density of the blades increases proportionally with the distance from the ladar.

4. Experimental Results

In this section we present some experiments on real range data, collected by our urban robot at La Canada, near JPL, in environments that include grass with some rocks and trees. The rotation rate of the ladar was set to 5Hz and the angular sampling period to 0.7° .

Figure 4 shows the spatial distribution of range measurements with the robot placed in front of a patch of grass. The grass is visible in the upper half-plane; the dots in the lower plane correspond to soil, which was visible because the rotation axis was not perfectly vertical. Two histograms of range are presented in Figure 5, covering respectively a narrow and a large angular sector of the grassy patch. As expected, the first histogram has a shape very similar to the exponential curve of (1). The best-fitting exponential (superimposed on the histogram) has mean $\mu = 1/\lambda d$ with $\lambda d = 3.7/\text{m}$. The second histogram has a much heavier tail, as predicted by the discussion in Section 3.2. Note in passing that the histograms have been computed over a number of revolutions while the robot stood still. A light breeze is sufficient to keep the grass blades in constant shaking motion, enabling a good statistical sampling. Another reason for scan-to-scan differences is that in each revolution the measurements may not be taken at identical angular positions.

Figure 6 shows a situation with some grass and two rocks (one of which partially occluded by grass). The histograms, computed within the angular sector corresponding to three different situations (grass, rock, rock behind the grass), match the densities predicted in Section 3 rather well. In particular, the histogram corresponding to the rock behind the grass is clearly bimodal (see also Figure 8).

A classifier of grass versus non-grass may be based on the statistical properties of the range discussed above. We implemented a very simple and fast algorithm, based the local estimation of the variance and of the skewness of the range distribution. A suitable threshold set on the variance may allow us to classify rather robustly grass from an obstacle with an exposed smooth surface. However, this technique will fail when the obstacle is partially occluded by grass, because the range measurements will still have a high variability. To deal with this situation, we may exploit the fact that the skewness is lower than in the case of homogeneous grass (see Section 3.3). Thus, our improved statistical test classifies a point as grass if the local variance is above a threshold t_1 and the local skewness is above a threshold t_2 .

An important issue here is the choice of the sample size for the computation of the variance. A small sample size causes the variance to vary widely within the same terrain cover class. A large sample size covers a broad area,

and therefore reduces the spatial resolution of the estimation. It was found by extensive testing that a sample size of 9 for the variance and of 29 for the skewness represents a good compromise between resolution and stability. Figure 9 shows the classification results using our algorithm in two different situations.

5. Conclusions

We presented an analysis of the statistics of range measurements in a vegetated environment, and showed its use in the design of a classifier of grass versus obstacles. Our technique has given good results in real-world experiments, even when obstacles were partially occluded by grass. Future work will extend our analysis to more complex situations (involving, for example, discriminating patches of visible soil), as well as to the case of two-dimensional range data.

Acknowledgments

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, and was sponsored by the Defense Advanced Research Projects Agency under the Mobile Autonomous Robot Software program and under the Tactical Mobile Robotics program. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

References

- [1] P. Bellutta, R. Manduchi, L. Matthies, K. Owens, A. Rankin, "Terrain Perception for DEMO III", *IEEE Intelligent Vehicles Symposium 2000*, Dearborn, MI, October 2000
- [2] C. Bergh, B. Kennedy, L. Matthies, A. Johnson, "A compact and low power two-axis scanning laser rangefinder for mobile robots", *Seventh Mechatronics Forum International Conference*, Atlanta, Georgia, September 2000.
- [3] J. Huang, A.B. Lee, D. Mumford, "Statistics of range images", *IEEE Conference on Computer Vision and Pattern Recognition*, Hilton Head, June 2000.
- [4] M. Kendall, A. Stuart, *The advanced theory of statistics*, Macmillan Publishing, 1977.
- [5] T. Leung, J. Malik, "On perpendicular texture or: Why do we see more flowers in the distance?", *IEEE Conference on Computer Vision and Pattern Recognition*, San Juan, Puerto Rico, June 1997.
- [6] L. Matthies, Y. Xiong, R. Hogg, D. Zhu, A. Rankin, B. Kennedy, "A portable, autonomous, urban reconnaissance robot", *International Conference on Intelligent Autonomous Systems*, Venice, Italy, July 2000.
- [7] L. Matthies, A. Kelly, T. Litwin, G. Tharp, "Obstacle detection for unmanned ground vehicles: a progress report", *Robotics Research: Proceedings of the 7th International Symposium*, 1996.
- [8] C. Nikias, A. Petropulu, *Higher-order spectral analysis*, Prentice-Hall, 1993.

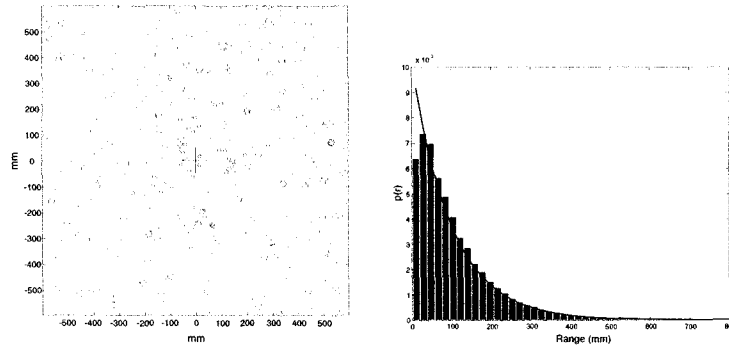


Figure 1. Left: Simulated distribution of grass ($\lambda = 500\text{m}^2$, $d = 20\text{mm}$). The laser is placed in the middle; the acquisition angular period is 0.5° . The circles filled in red represents blades of grass hit by a laser ray. Right: Normalized histogram of range over 500 trials.




Figure 2. Left: Simulated distribution of grass ($\lambda = 85/\text{m}^2$, $d = 30\text{mm}$). The laser is placed in the lower left corner, at a distance of 200 mm from the patch of grass. Right: Normalized histogram of range over 500 trials.




Figure 3. Left: Simulated distribution of grass ($\lambda = 300/\text{m}^2$, $d = 10\text{mm}$). The blue circles represent an impenetrable surface. The laser is placed in the lower left corner. Right: Normalized histogram of range over 300 trials.

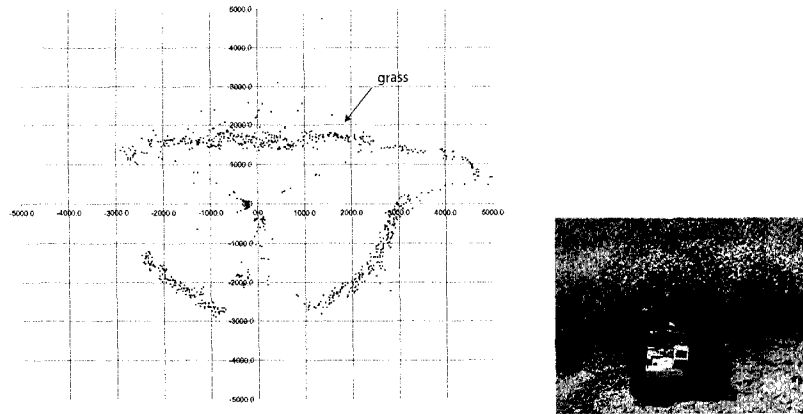


Figure 4. Spatial distribution of range measurements around the robot. In this figure as well as in the following ones, the axis ticks represent millimeters. The cluster of points at the immediate left of the center is actually a part of the robot within the field of view of the ladar.




Figure 5. Normalized histogram of range measured over 30 revolution at angles between 60° and 120° (left) and between 5° and 140° (right) for the case of Figure 4.




Figure 6. A situation with two rocks and grass.




Figure 7. Normalized histogram of range measured at angles between -30° and 30° (left), between -40° and -60° (center) and between 35° and 50° (right) for the case of Figure 6.




Figure 8. A situation with a rock behind the grass (left) and the normalized histogram of range measured over 3 revolution at angles between -45° and -70° (right).




Figure 9. Classification results using local measurements of range variance and skewness for two different instances. The red enlarged dots have been identified by the algorithms as non-grass.